

OPERATIONAL DATA STORE

ABSTRACT OF THE DISCLOSURE

An operational data store consists of an insert table for storing new data and a history table, partitioned by range and further sub-partitioned, for storing historical data.

- 5 Transfer logic periodically transfers new data from the insert table to the history table. The transfer logic includes a secondary table and fill logic for filling the secondary table with selected data from the insert table. Secondary transfer logic transfers the secondary table into the history table, such that the selected data is transferred into the history table. Indexing logic applies the history table indexing scheme to the secondary table.
- 10 Table logic creates a new partition the history table, for swapping with the secondary table, by exchanging respective pointers. A query engine may apply a database query to both the history table and the insert table, so that all data is available. An aggregator accumulates new data into an aggregation buffer. The accumulated data are batched and transferred into the insert table with a single database access. A throttler throttles
- 15 transactions of different classes and types independently to achieve a desired level of service. The system can be configured to execute in a plurality of processor nodes configured as a processor cluster, wherein distinct database server instances are associated with distinct processor nodes of the processor cluster.